University of California
Lawrence Livermore National Laboratory

High-Performance Global Climate Modeling

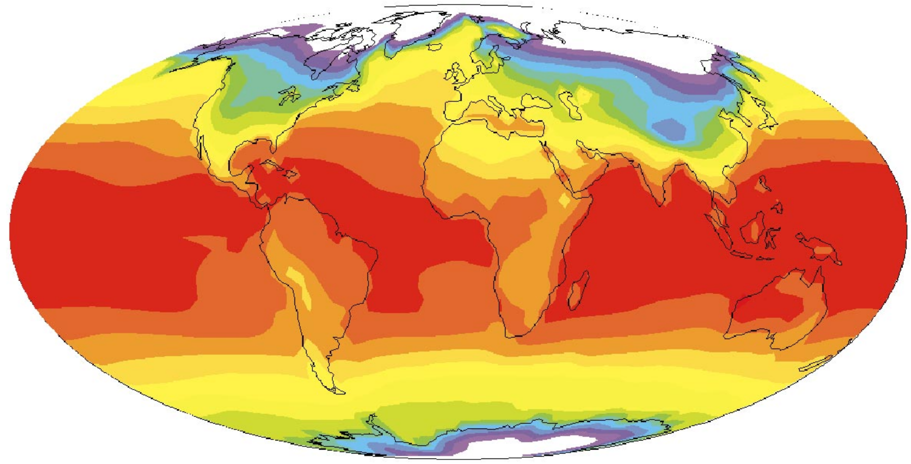

Mean surface temperature during the northern hemisphere's winter months, as computed by our atmospheric general circulation model.

Mission

We are developing an advanced generation of parallel climate models together with a computational framework that will provide a comprehensive climate systems modeling capability on high-performance computing systems.

Impact

A better understanding of the complex interplay among the many processes that affect the climate will help us estimate more accurately the impact of greenhouse gases and other products of human activity.

As part of our long-range plan to develop a comprehensive climate systems modeling capability, we have developed a Climate Systems Modeling Framework (CSMF) for high-performance computing systems, designed to schedule and couple multiple physics simulation packages in a flexible and transportable manner. We have also taken leading oceanic and atmospheric general circulation models, enhanced them, and recast them in parallel form. Parallelism is achieved through both domain decomposition and process-level concurrency, with data transfer and synchronization accomplished through message-passing. Both machine transportability and architecture-dependent optimization are handled through libraries and conditional compile directives. Scientific studies using the separate oceanic and atmospheric models along with preliminary coupled ocean / atmosphere calculations are being performed on a number of high-performance platforms, including the CRAY T3D, IBM SP2, Meiko CS-2, and Intel Paragon.

Atmospheric and Oceanic Models

Both the atmospheric and oceanic components use a finite-difference methodology on a three-dimensional Eulerian mesh. The atmospheric model was originally developed at UCLA. It consists of a hydrodynamics module, which is based on a simplification of the Navier Stokes equations, and a column physics module, which consists of parameterizations of a number of processes important to the simulation of the atmosphere. The hydrodynamics finite-difference mesh is staggered both horizontally and vertically. This facilitates the conservation of key physical quantities, which is critical for long-term climate simulations. The physics module addresses long-wave and short-wave radiation, cumulus convection, and other important processes.

The oceanic model comes from the National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory. It includes detailed bathymetry and coastline boundaries. We have extended the code to allow a more general interface between the ocean and atmosphere. The calculation is divided into two phases, known as baroclinic and barotropic. The baroclinic phase is similar to the hydrodynamic time-advance in the atmospheric code. The barotropic phase consists of updating the vertically averaged velocity components; the solution procedures are vastly different depending on the ocean / atmosphere interface. We have extended the oceanic code to model polar sea ice through coupling to a dynamical thermodynamic sea ice model.

Climate System Modeling Framework

The CSMF binds together the atmospheric circulation, oceanic circulation, and sea ice models discussed above, as well as models of atmospheric chemistry land surface processes and ocean biogeochemistry. Package execution is managed by a high-level scheduler based on the dataflow concept. That is, a package is time-advanced until it is due to receive data from a different package.

The coupling aspects are nontrivial both from the physics and computational standpoints. In addition to identifying the data to be exchanged, one must assure conservation of key physical quantities. Also, one must take into account the different mesh sizes and domain decompositions of the various packages. We have addressed some of these issues by creating separate coupling modules, in particular for the ocean / atmosphere interface.

Parallelization Strategy

We accomplish parallelism within the oceanic and atmospheric models by partitioning the global spatial domain into a collection of subdomains and time-advancing the various subdomain solutions concurrently. Processors are assigned to subdomains in a deterministic manner, and variables local to a given subdomain are stored on the memory of the processor assigned to that subdomain. Data are transmitted between subdomains in the form of messages. Both packages use latitude, longitude, and a vertical coordinate as independent variables. A two-dimensional latitude / longitude domain decomposition is implemented, whereby each subdomain consists of a number of contiguous columns having full vertical extent appropriate to that package.

Process-level concurrency is attained by partitioning the machine among physical packages. The number of processors per package must be chosen judiciously so that the computational load is roughly balanced. At present this is a static procedure.

Simulations of the Atmosphere

A standard means of validating models of atmospheric general circulation is provided by the Atmospheric Model Intercomparison Project (AMIP). The test problem is defined by simulation of the ten-year period from 1979 to 1989, with prescribed sea surface temperatures. Over thirty atmospheric codes worldwide have been compared using this test problem. Our acquired parallel processing technology has enabled us to execute an ensemble of twenty such realizations and to calculate the variability of the model atmosphere. This provides an estimate as to the natural variability of the atmosphere, without which we could not assess human-induced climate change. We believe this ensemble to be the largest yet performed by any one group.

The figure on the front page shows the mean surface temperature during the northern hemisphere winter months. A calculation of the average standard deviation of the ensemble mean surface temperature (not shown) indicates that the variation is largest over regions of sea ice and snow-covered land, suggesting a radiative feedback mechanism.

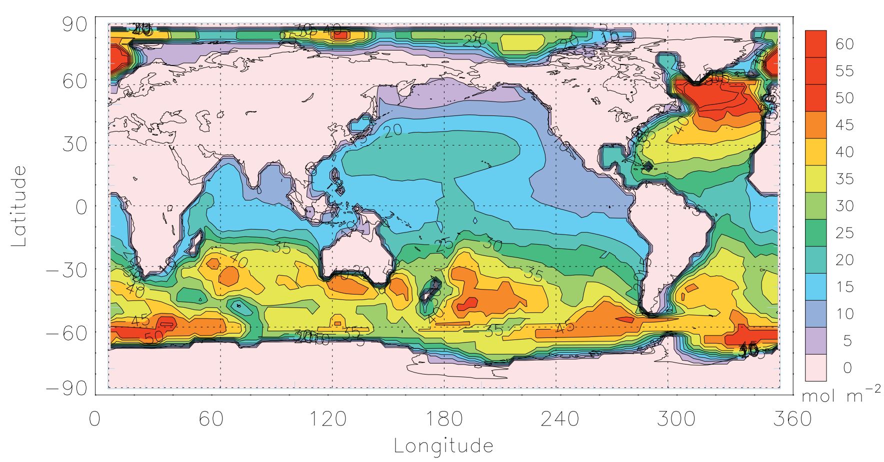

Column inventory of anthropogenic carbon dioxide, as computed by our coupled ocean circulation/biogeochemistry model.

Simulations of the Ocean

We have been studying the carbon cycle using the coupled oceanic circulation biogeochemistry model. This combination provides the most sophisticated carbon isotope treatment in an ocean code to date. Such comprehensive computations are made possible by our parallel processing capability. One use of this coupled model is to estimate the degree of uptake of anthropogenic (human-produced) carbon dioxide by the ocean. The figure below shows a column inventory (i.e., the amount contained in a vertical column) of anthropogenic carbon dioxide as a function of latitude and longitude. This calculation takes into account estimates of atmospheric carbon dioxide emanating from human activities. This computation led to our discovering a heretofore unobserved flux of carbon-14 from the ocean to the atmosphere induced by fossil fuel burning.

Future Directions

We are now using the CRAY T3D for scientific study of the ocean and atmosphere, and preliminary coupled ocean / atmosphere calculations are being carried out as well. Some optimization remains to be done, particularly in the area of load balancing. Our work to date suggests that it should be possible to develop an advanced coupled climate simulation model with performance scalable to the teraflops range.

For more information about high-performance global climate modeling, contact William P. Dannevik, 510-422-3132, dannevik1@llnl.gov; or Arthur A. Mirin, 510-422-4020, mirin@llnl.gov.

Work performed under the auspices of the U.S. DOE by LLNL under contract No. W-7405-ENG-48